



Lake Partner Onboarding Guide

AWS CloudTrail



Version 1.0

Copyright © 2024 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

AWS CloudTrail: Lake Partner Onboarding Guide

Copyright © 2024 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

Onboarding partner events to AWS CloudTrail Lake	1
How integrations with CloudTrail Lake add value	1
Terminology	2
How partner integration works	3
Prerequisites	4
Onboard to AWS CloudTrail Lake	4
Step 1: Partner registration	4
Step 2: Build the integration	5
Best practices and quotas	5
CloudTrail Lake event schema	6
Learning resources	14

Onboarding partner events to AWS CloudTrail Lake

AWS CloudTrail Lake logs activities across all AWS accounts and Regions, and across a customer's entire IT infrastructure. Customers can configure CloudTrail Lake to log events from any source, immutably store the events for auditing and compliance, and use standard, SQL-based queries to filter and analyze their event logs. CloudTrail Lake accepts activity logs from AWS Partner Network partner solutions, offering customers a comprehensive view of their activity information in the CloudTrail Lake console, or by using API commands.

This guide is for AWS Partners who want to explore creating an integration with CloudTrail Lake, or want to know how to onboard their applications and solutions to CloudTrail Lake, and let their customers integrate their activity events into CloudTrail Lake.

Topics

- [How integrations with CloudTrail Lake add value](#)
- [Terminology](#)
- [How partner integration works](#)
- [Prerequisites](#)
- [Onboard to AWS CloudTrail Lake](#)
- [Learning resources](#)

How integrations with CloudTrail Lake add value

As an AWS Partner, an integration with CloudTrail Lake can add value for you in the following ways:

- **Modern, consolidated solution for audit logging for your customers:** Today, audit and security professionals get trusted records of AWS activity from CloudTrail. CloudTrail users need a similar experience for other application audit information, regardless of the source. Partner integrations centralize audit logging, and extend CloudTrail benefits, such as immutable storage for 7 years and a query interface for analysis, to partner solutions, simplifying audit and compliance processes for our common customers.
- **Discovery for partners:** CloudTrail promotes partners with integrations in CloudTrail console, including links to partner AWS Marketplace listings.

Terminology

Event data store

CloudTrail Lake lets you run SQL-based queries on your events. Events are aggregated into event data stores, which are immutable collections of events based on criteria that you select by applying [advanced event selectors](#). You can keep the event data in an event data store for up to 3,653 days (about 10 years) if you choose the **One-year extendable retention pricing** option, or up to 2,557 days (about 7 years) if you choose the **Seven-year retention pricing** option. The selectors that you apply to an event data store control which events persist and are available for you to query. For non-AWS sources, including partner sources, customers create an event data store to log activity events using the CloudTrail console or API. The event type in the console is **Events from integrations**. In the API, the `eventCategory` value is `ActivityAuditLog`.

Channel

A partner-specific resource that AWS customers create as part of the integration process in CloudTrail Lake. Channels let customers map event sources to destinations. Channels for onboarded partner events have the partner solution set as the source, and event data stores to which customers want to deliver partner events set as destinations. To finish the integration process, customers provide the partner with an Amazon Resource Name (ARN) of the channel. The partner solution uses the channel to send events to CloudTrail Lake.

Resource policy

A permissions policy that is attached to the channel resource and identifies who has access to the channel.

Direct integration

CloudTrail supports two integration types: direct and solution. With a direct integration, the partner calls the `PutAuditEvents` API to deliver events to the event data store for the customer's AWS account.

Solution integration

CloudTrail supports two integration types: direct and solution. With a solution integration, the application runs in the customer's AWS account and the application calls the `PutAuditEvents` API to deliver events to the event data store for the customer's AWS account.

6. CloudTrail Lake checks the resource policy to verify that the partner's permissions are valid. If the partner's permissions are valid, CloudTrail Lake ingests the activity events.

Prerequisites

The following are requirements for performing tasks in this guide.

- You must be at least an [AWS Select Tier Partner](#). For more information about how to become an AWS Select Tier partner, see [Join AWS Partner Network](#).

Onboard to AWS CloudTrail Lake

Complete tasks in this section to onboard your partner application to CloudTrail Lake.

Topics

- [Step 1: Partner registration](#)
- [Step 2: Build the integration](#)
- [Best practices and quotas](#)
- [CloudTrail Lake event schema](#)

Step 1: Partner registration

To get started, register as an AWS Partner in the AWS Partner Network.

Be sure to meet the requirements of partner intake forms. The partner CloudTrail Lake intake forms collect information that the AWS Partner Network uses to create your partner product profile. This profile gives the CloudTrail team information that we add to your partner provider description that is displayed in the CloudTrail console. Your profile also includes information that CloudTrail uses to confirm the integrity of the event source as CloudTrail Lake receives events the from a partner application.

1. Get started by [joining the AWS Partner Network](#), and informing your AWS Partner Network team that you want to become a partner with CloudTrail Lake.
2. Get onboarding materials—including partner onboarding forms and the CloudTrail event schema—from the AWS Partner Network team.

3. Finish partner onboarding forms, and share with your AWS Partner Network team. You might not yet have all required details. If you have questions, contact your AWS Partner Network team.

Step 2: Build the integration

Build the integration that is required to send event logs to CloudTrail Lake.

1. Review the [CloudTrail integrations event schema](#) in this guide. The CloudTrail event schema provides a consistent way to log activity events for audit needs. This eliminates the need for time-consuming data standardization efforts before a cross-source analysis. CloudTrail Lake cannot accept events that do not follow the prescribed schema.
2. Determine the events that you want to send. CloudTrail Lake only accepts activity events, or events that help customers understand who did what, and when. Typically, partners have existing mechanisms to provide their customers access to activity logs. The schema mapping exercise helps you exclude non-activity events. Contact your AWS Partner Network team if you need help narrowing down event types.
3. Build your integration architecture to send activity events to CloudTrail Lake. This includes offering a setup framework (GUI is preferred) and documentation for customers to enable your partner application to send events to CloudTrail Lake. A partner customer must share a CloudTrail channel Amazon Resource Number (ARN) with the partner as part of the integration process.
 - a. To send events to CloudTrail Lake, the partner calls the [PutAuditEvents API](#), specifying the channel ARN provided by the customer. If the channel's resource policy includes an external ID, you must also pass the external ID when you call PutAuditEvents.
 - b. The partner checks transfer results for failures, and tries to resend failed events by calling the PutAuditEvents API again.

Best practices and quotas

As you integrate partner solution events, be aware of the following best practices, quotas, and limitations.

- **Schema mapping:** Be sure that you have the key required fields included in the `eventData` block. Missing required fields results in errors. For information about required fields, see [CloudTrail Lake event schema](#)

You can add event fields that do not map to the schema to the `additionalEventData` field. Some partners use this field to include the entire, raw event.

- **Batching events:** When you call the `PutAuditEvents` API, you can batch up to 100 events in a single API call, as long as each event is not greater than 256 kB in size, and the total size of all events is less than 1 MB. For more information about quotas in CloudTrail, see [Quotas in AWS CloudTrail](#) in the *AWS CloudTrail User Guide*.

CloudTrail Lake event schema

The following table describes the required and optional schema elements that match those in CloudTrail event records. The contents of `eventData` are provided by customer events; other fields are provided by CloudTrail after customer events are ingested.

- [Fields that are provided by CloudTrail after ingestion](#)
- [Fields that are provided by your events](#)

The following fields are provided by CloudTrail after ingestion:

Field name	Input type	Requirement	Description
<code>eventVersion</code>	string	Required	The event version.
<code>eventCategory</code>	string	Required	The event category. For non-AWS events, the value is <code>ActivityAuditLog</code> .
<code>eventType</code>	string	Required	The event type. For non-AWS events, the valid value is <code>ActivityLog</code> .

Field name	Input type	Requirement	Description
eventID	string	Required	A unique ID for an event.
eventTime	string	Required	Event timestamp , in yyyy-MM-DDTHH:mm:ss format, in Universal Coordinated Time (UTC).
awsRegion	string	Required	The AWS Region where the PutAuditEvents call was made.
recipientAccountId	string	Required	Represents the account ID that received this event. CloudTrail populates this field by calculating it from event payload.
addendum	-	Optional	Shows information about why event processing was delayed. If information was missing from an existing event, the addendum block includes the missing information and a reason for why it was missing.

Field name	Input type	Requirement	Description
• reason	string	Optional	The reason that the event or some of its contents were missing.
• updatedFields	string	Optional	The event record fields that are updated by the addendum. This is only provided if the reason is UPDATED_DATA .
• originalUID	string	Optional	The original event UID from the source. This is only provided if the reason is UPDATED_DATA .
• originalEventID	string	Optional	The original event ID. This is only provided if the reason is UPDATED_DATA .
metadata	-	Required	Information about the channel that the event used.
• ingestionTime	string	Required	The timestamp when the event was processed, in yyyy-MM-DDTHH:mm:ss format, in Universal Coordinated Time (UTC).

Field name	Input type	Requirement	Description
<ul style="list-style-type: none"> channelARN 	string	Required	The ARN of the channel that the event used.

The following fields are provided by customer events:

Field name	Input type	Requirement	Description
eventData	-	Required	The audit data sent to CloudTrail in a PutAuditEvents call.
<ul style="list-style-type: none"> version 	string	Required	<p>The version of the event from its source.</p> <p>Length constraints: Maximum length of 256.</p>
<ul style="list-style-type: none"> <ul style="list-style-type: none"> userIdentity 	-	Required	Information about the user who made a request.
<ul style="list-style-type: none"> <ul style="list-style-type: none"> <ul style="list-style-type: none"> type 	string	Required	<p>The type of user identity.</p> <p>Length constraints: Maximum length of 128.</p>
<ul style="list-style-type: none"> <ul style="list-style-type: none"> <ul style="list-style-type: none"> principalId 	string	Required	A unique identifier for the actor of the event.

Field name	Input type	Requirement	Description
			Length constraints: Maximum length of 1024.
<ul style="list-style-type: none"> • details 	JSON object	Optional	Additional information about the identity.
<ul style="list-style-type: none"> • userAgent 	string	Optional	<p>The agent through which the request was made.</p> <p>Length constraints: Maximum length of 1024.</p>
<ul style="list-style-type: none"> • eventSource 	string	Required	<p>This is the partner event source, or the custom application about which events are logged.</p> <p>Length constraints: Maximum length of 1024.</p>
<ul style="list-style-type: none"> • eventName 	string	Required	<p>The requested action, one of the actions in the API for the source service or application.</p> <p>Length constraints: Maximum length of 1024.</p>

Field name	Input type	Requirement	Description
<ul style="list-style-type: none"> eventTime 	string	Required	Event timestamp , in yyyy-MM-DDTHH:mm:ss format, in Universal Coordinated Time (UTC).
<ul style="list-style-type: none"> UID 	string	Required	<p>The UID value that identifies the request. The service or application that is called generates this value.</p> <p>Length constraints: Maximum length of 1024.</p>
<ul style="list-style-type: none"> requestParameters 	JSON object	Optional	The parameters, if any, that were sent with the request. This field has a maximum size of 100 kB, and content exceeding the limit is rejected.
<ul style="list-style-type: none"> responseElements 	JSON object	Optional	The response element for actions that make changes (create, update, or delete actions). This field has a maximum size of 100 kB, and content exceeding the limit is rejected.

Field name	Input type	Requirement	Description
<ul style="list-style-type: none">errorCode	string	Optional	<p>A string representing an error for the event.</p> <p>Length constraints: Maximum length of 256.</p>
<ul style="list-style-type: none">errorMessage	string	Optional	<p>The description of the error.</p> <p>Length constraints: Maximum length of 256.</p>
<ul style="list-style-type: none">sourceIPAddress	string	Optional	<p>The IP address from which the request was made. Both IPv4 and IPv6 addresses are accepted.</p>
<ul style="list-style-type: none">recipientAccountId	string	Required	<p>Represents the account ID that received this event. The account ID must be the same as the AWS account ID that owns the channel.</p>

Field name	Input type	Requirement	Description
• additionalEventData	JSON object	Optional	Additional data about the event that was not part of the request or response. This field has a maximum size of 28 kB, and content exceeding that limit is rejected.

The following example shows the hierarchy of schema elements that match those in CloudTrail event records.

```
{
  "eventVersion": String,
  "eventCategory": String,
  "eventType": String,
  "eventID": String,
  "eventTime": String,
  "awsRegion": String,
  "recipientAccountId": String,
  "addendum": {
    "reason": String,
    "updatedFields": String,
    "originalUID": String,
    "originalEventID": String
  },
  "metadata" : {
    "ingestionTime": String,
    "channelARN": String
  },
  "eventData": {
    "version": String,
    "userIdentity": {
      "type": String,
      "principalId": String,
      "details": {
        JSON
      }
    }
  }
}
```



```
    }
  },
  "userAgent": String,
  "eventSource": String,
  "eventName": String,
  "eventTime": String,
  "UID": String,
  "requestParameters": {
    JSON
  },
  "responseElements": {
    JSON
  },
  "errorCode": String,
  "errorMessage": String,
  "sourceIPAddress": String,
  "recipientAccountId": String,
  "additionalEventData": {
    JSON
  }
}
```

Learning resources

The following resources can help you get a better understanding of what CloudTrail Lake is and how AWS customers use it. We encourage you to try CloudTrail in one of your AWS accounts, and get more experience using the service.

- [Modernize Your Audit Log Management Using CloudTrail Lake](#) (YouTube video)
- [Log Activity Events from Non-AWS Sources in AWS CloudTrail Lake](#) (YouTube video)
- [Get visibility into the activity logs for your workforce and customer identities](#) (AWS blog)
- [Using AWS CloudTrail Lake to identify older TLS connections to AWS service endpoints](#) (AWS blog)
- [How Arctic Wolf uses AWS CloudTrail Lake to Simplify Security and Operations](#) (AWS blog)
- [Working with CloudTrail Lake](#) in the *AWS CloudTrail User Guide*
- [AWS CloudTrail API Reference](#)
- [AWS CloudTrail Data API Reference](#)